

OpenAI starts GPT-5.6 in limited preview at the US government's request

Source copy kept by wheresthe.ai for the story "OpenAI starts GPT-5.6 in limited preview at the US government's request" (<https://wheresthe.ai/d/openai-gpt-5-6-limited-preview-2026-06/>).

Source: <https://openai.com/index/previewing-gpt-5-6-sol/>

Text saved 18 Oct 2026 (the page text through a reader, because the site refused a direct fetch).

SHA-256 of the saved text: 492818fd65840e97dba84e2b08192a0db566eee01bcb0535a7933f9b14b12dc3

[Skip to main content](#)

[* Research](#) [* Products](#) [* Business](#) [* Developers](#) [* Company](#) [* Foundation](#)(opens in a new window)

[Log in](#)[Try ChatGPT](#)(opens in a new window)

[* Research](#) [* Products](#) [* Business](#) [* Developers](#) [* Company](#) [* Foundation](#)(opens in a new window)

Previewing GPT-5.6 Sol: a next-generation model | OpenAI

0 0 . 0 . . . 0 . . . 0 . 0 0 0 0 . 0 . . . 0 . . . 0 0 . . . 0 . 0 0 0 0 . . . 0 . 0 . . . 0 0 . . . 0 .
0 0 0 . 0 0 . . . 0 0 . . . 0 0 0 0 0 . 0 0 . . . 0 . . 0 0 . 0 . 0 0 0 0 0
. . . 0 . . . 0 0 0 0 0 . . . 0 0 0 . . . 0 0 . 0 0 . 0 . 0 . . 0 0 . 0
. 0 . . 0 . . . 0 0

655655555555 5555566655555.665 555556565.55555655555665 566656655555.565665556565.65
555.55555655556565.56656.555556 5555.555555.55565655555.5665555655
56.6.565.5.66..55655665.55555555.556 56.566555.556556556.665556556565665.5.55
55555.55555655555.55565.565.66555555.5.5 65.555566555...565555..5.555.565555.56.655
556.5555656556.56555.55556556555555565665565 5556.555555555556555.5566.5555.556656.655556
6.55555655556565565.55.555565555555555.66555 .66555.5.6555555655566555555.665556.5555.565566
5555.65555655.65.5.5565.55.655..55656.55555555 5665.5555565656.55556.5..6656565555656555.65555
56.55655566555566565...5.5556.56.5655.55555655555
65555555555655555555555.555565565556.55.55565556
5555.655555565655555556555.56665556566556.5556.5655
.5665.5.555555556555656655656556.665556.555555
56656.565565555.5556.55555655555.5.55655.65.66.56
56.55.65555555555555555.565..5566655555555.55555
5555.556555665.555555.55565555.555665.55556655
.6655656566.55666555656565656.5..55.6.65556655
55.565.5566655555566.555555555566..5555556..6 55.55556655.56566.6565655.555.555655.5655555555
665555555.555556665.656655565565.5556556556556 655.5555555555.5.55.6565566.5566565.55555.55
56555565.556.5666666555.6566555565556655.5565 5.55565655555555555655..5.5.65555655565555
56655665556565655555.66555556555565566655 .6655665.665555.556556556656655555.555
55555565555556556555555.5.55.66656555. 5565555556665.5.655566556.665565555
6.66566555555..55555555555655555 556656555555.555555555555555.5556
56..5556665655.5565555555556 .55555655655..666555655 6565656.5555555555 56555555665 565 5 6 5 5
5 5 5 6 5 5 . 5 5 5 6 5 5 5 6 5 5 5 5 6 . 5 5 5 5 5 5 6 5 6 5 5 . 5 5 5 5 5 . 5 . 5 5 6 5 5 5 6 6 5 . 5 6 5 5 6 5 6 5
5 5 5 5 . 6 5 6 5 5.55 5 . 6 6 5 5 5 6 5 . 5 5 5 5 6 . 6 5 5 . 6 6 5 5 5 5 5 6 5 5 6 5 5 5 5 6 5 5 5 5 . 5 6 5 6 5 6 6 5 5 5
6 5 5 6 5 . 5 5 5 5 5 5 5 5 5 5 5 6 5 6 . 5 . 5 6 5 5 . 5 5 . . 6 5 . 5 . 5 . 6 5 5 5 5 6 5 6 5 5 5 6 5 5 5 5 5 5 5 6 6
6 5 6 5 5 . 5 5 6 5 5 5 5 5 . 6 5 5 5 5 5 5 5 5 . 5 5 5 5 6 5 5 . 5 5 5 5 5 5 5 . 5 5 5 5 . 5 ..6 5 5 5 5 5 5 5 5 5 6 5 5
5 5 5 5 6 6 5 5 5 . 5 6 5 5 5 5 6 . 5 6 5 5 . 5 5 5 6 6 5 5 5 6 6 . 6 5 5 6 5 5 5 5 5 5 6 6 5 5 5 5 5 6.55 5 6 6 5 5 5 6
5 5 5 5 5 . . 5 . 6 5 5 5 5 6 5 5 . 5 5 5 6 5 5 5 6 . . 5 5 5 5 5 5 5 5 6 5 5 5 6 5 5 5 5 6 . 5 5 6 5 6 5 5 5 . .
6 6 6 . 5 5 5 5 5 5 5 5 5 5 6 6 . 5 5 5 5 6 6 5 6 5 5 5 5 . . 5 5 5 5 6 . 5 5 5 5 6 6 . 5 5 6 6 . 5 5 5 6 5 6 5 5 6 6
5 6 5 5 6 6 5 5 6 6 5 6 5 6 5 6 5 6 5 6 . 5 5 6 6 5 . 5 5 6 5 6 5 5 5 . 5 6 5 5 6 6 5 5 5 5 5 5 5 5 5 6 6 6 5
5 5 6 . 6 5 . 5 5 5 5 5 5 5 5 5 6 5 . 5 6 5 5 5 5 5 6 5 6 5 5 . 5 6 6 5 . 6 5 . 5 . 5 6 5 5 6 5 . 5 5 5 6 5 5 6 5 6 5 5 5 5
5 . 5 5 5 6 . 5 5 5 5 5 6 5 6 5 . 5 5 5 5 5 6 5 6 . . 5 5 6 . 5 5 5 5 5 5 . 5 5 6 5 5 6 5 . 5 5 5 5 5 5 6 5 6 5 5 5 5
6 5 . 5 . 6 5 . 5 . 5 5 5 5 . 5 5 5 . 5 6 5 6 5 5 6 . 6 5 5 6 5 5 5 6 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 5 . 5

5 555 6 56 6 5 5 6665 6 6 55 6 5 . 6 55 . 5 6 5 6 6 5 5 5 5 6 . 5 5 5 6 55 5 6 65 5 5 5 5 6 5 55 6 5 6 66 5 5 5 6 6 5 .
55 5 5 6 . 5 5 6 5 5 6 6 6 5 55 6 5 5 5 . 5 5 . . 5 6 . 6 5 5 5 . . 5 . 5 55 5 5 5 5 5 6 5 5 6 5 5 5 5 65 5 5 6 5 . 6 6 5
555 5 66 5 5 5 . . 5 6 5 . 5 5 5 6 5 6 5 5 6 5 65 5 55 55 5 5 556 5 6 556 55 5 6 5 5 6 5 5 6 5 5 5 56 5 6 5 5 5 5 6
5 5 . 55 5 5 5 655 5 5 5 6 5 5 . 5 5 5 6 5 5 6 6 5 5 6 5 555 6 5 6 5 5 5 5 6 5 5 6 5 5 5 6 5 5 . 5 5 6 5 5 5 5 5
5 . 5 5 5 6 5 55 5 6 5 5 . 5 5 5 5 5 5 56 55 5 6 5 5 6 6 5 6 55 . 5 5 5 6 5 6 5 6 5 .

5 6 . 5 5 55 6 5 5 5 5 5 6 5 5 5 5 5 5 6 5 6 55 5 5 6 5 5 5 55 5 5 5 5 5 . 56 6 5 5 5 6 6 . 6 6 6 6 5 5 . 55 55 5 5 5 5 5
5 6 5 . 5 . 5 5 . 5 5 6 5 5 5 5 6 . 5 6 . 5 6 5 . 5 5 55 5 5 6 5 55 6 56 5 5 56 5 5 5 5 5 5 . 6 6 5 6 . 5 5 5 5 55 55 555
5 5 6 5 . 5 65 . 5 5 55 56 5 5 5 . 5 5 5 55 5 5 . 6 5 5 . 5 5 65 5 5 6 5 6 65 6 6 5 . 55 5 5 6 55 . . 5 5 5 6 66 5 5 5 5 55
565 5 6 . 5 5 5 5 6 6 5 66 5 5 5 5 5 . . 5 56 6 5 65 5 5 5 5 5 5 5 5 . 5 6 56 6 . 6 5 5 55 5 5 5 5 55 . 6 6 5 5 5 5 5 5 6
5 5 6 6 66 5 6 5 56 6 6 5 5 . 5 . 5 . 6 . 5 . 6 . . 5 . 6 5 5 55 55 5 5 . 5 5 5 . 55 5 5 5 5 5 6 5 5 . 5 5 55 5 5 5 6 5 5 5
6 5 . 5 5 5 6 . 56 5 5 6 5 56 5 . 5 5 6 5 5 5 5 5 5 . 5 . 6555 56 . 656 5 6 . 5 . 5 . 5 5 5 6 6 5 6 5 565 . 6 5 5 555 5 65
5 5 6 65 56 . 5 . 5 6 5 55 5 6 5 5 6 6 6 6 5 . 5 6 55 5 5 5 5 . 65 6 5 5 5 5 55 5 5 . . 5 . 6 5 5 5 . . 55 5 5 . 55 55 5 5 5 6
5 55 . 5 56 5 5 6 5 65 655 65 5 5 5 5 . 5 66 5 . 5 6 5 5 55 5 5 5 55 . 6 6 56 6 5 5 5 5 6 5 6 5 5 5 6

5 6 5 6 6 6 5 5 5 5 5 5 5 6 5 . 5 5 5 5 5 5 5 . 5 5 . 5 5 5 5 5 6 5 6 5 5 5 5 6 5 5 5 6 5 5 5 6 5 5 5 6 5 5 5 5 5 6
5 5 6 5 6 5 5 6 5 5 5 5 5 5 5 5 5 5 5 5 5 6 6 5 5 5 5 5 5 6 6 5 5 5 5 5 6 6 5 5 5 5 6 5 5 5 6 5 5 5 6 5 5 5 6
5 5 6 5 5 5 5 5 6 6 6 5 5 5 5 5 6 6 5 5 5 5

5

5

5.6...6.5..6..5..

5

5

5

5

5

5

565...65 6555.5..6..565 .5....566.5.555655 56.55.6555.5566656.5 5565.5556.5...6..656.6
.65.56.566555555.5656..5 656555565.65.6.5555566.556 66.565.55566565555.5.555..
5555556555.6565.6555556.6565 5...56.55555...5..655.55..56 5655.666665.5566.5655..56665
55655566665.65555555655555.. 5.66556.5...55.56.6.5.55.6 .6.5.65555..565.5555.55666..
565.5565..66556..56.555.655. 5.55556....5.6555.65..5..6 56566556.56655565556.6
55.5555555..6..565655.65 665.56565556555565.555 .555566.66..566..655 56.65655566.655... .65665556.6656
56.55555 65 . 6 5 5 5 . . 5 6 . 6 5 5 55 . 5 5 5 56 6 6 . . 6 6 6 55 5 . 56 5 5 6 5 5 . 6 5 5 66 5 . 55 55 5 . 5 . 6 . . 5 5 .
5 . . 66 . 5 6 5 55 6 6 5 . . 5 5 5 5 6 6 5 65 5 5 6 6 5 . 5 5 6 5 6 . 5 5 65 . 6 5 6 56 6 555 55 . . 6 6 . 5 6 . 5 6 . 6 . 5 . .
5 . 5 6 5 5 5 5 . 55 5 . 6 5 6 . 5 5 5 5 55 . 5 6 5 5 . 5 65 5 5 . . 5 5 6 6 . 6 5 . 6 56 5 6 5 55 5 66 56 .. 5 6 . . 5 5 5
55 5 6 5 . 65 . 56 55 . . 6 6 56 . 5 6 5 . 5 6 . 6 6 5 5 55 6 6 . 55 6 56 6 65 . 5 6 6 6 6 55 5 6 . . 5 5 5 . 5 . 6 6
. . 5 . 6 5 . 5 6 . 5 6 555 6 5 5 5 6 5 5 65 6 5 65 . 5 56 5 . 6 5 6 5 5 6 6 5 55 . 5 . 5 . 6 . 6 . . 5 5 5 . . 5 66 56 . . 5 . 5 .
. 5 . 5 56655 5 5 . . 5 5 6 65 55 555 . 5 5 6 . . . 6 5 555 6 55 5 . 5 5 6 5 6 5 . 6 5 6 . 6 6 . 5 5 5 55 5 . 6 . 56 6 5 56 . 65 5
6 5 6 556 6 . . 5 5 . 5 . 5 56 655 5 6 . 5 5 . 5 . 65 5 5 56 . 6 5 6 . 5 5 5 5 . 5 5 5 6 5 . 5 . 5555 . . 5 . 5 6 655 5
..6 65 55 5 5 6 5 5 5 6 6 . 6 6

.5.5 6 5..6 6 5 5.5 5 5 6 5 6.6 6.5 6...6..5 6 5 5.5 6.6 6 5 6

556565 56.5666.66 5555.6.6.... .5.5...56.556. 55656.55..655. 5....555555.55 .6.5.5566565.5 6.65.566.5..6.
556.555.5.6. 66.55.6555 6.5.55 5 6 6 66 5 6 5 5 . 5 5 5 6.5. 66 5 . 6.65 .6 5 6 5 5 5 5 . 5 6 . . 5 5 5 . 5 . . 5
. . 56 5 . 5 . 5 . . 6 . 5 565 6 5 6 5 5 . . 5 5 5 6 5 5 566 5 . 5 5 6 . 6 . 6 5 .. 6 6 6 5 . 5 . . 55 5 5 . . 56 5 5 6 . 65 5 .

5.6 5..5 5.6 5 5 5 5 5 5..5.5 6

June 26, 2026

ProductRelease

Previewing GPT-5.6 Sol: a next-generation model

Read the launch post

Listen to article 11:55

Audio 2

Share

Capabilities

* Capabilities * Stronger cyber capabilities with stronger safeguards * A layered safeguard stack * Improving robustness with automated red-teaming * Availability and pricing

* Capabilities * Stronger cyber capabilities with stronger safeguards * A layered safeguard stack * Improving robustness with automated red-teaming * Availability and pricing

We're beginning a limited preview of the GPT-5.6 series: Sol, our flagship model; Terra, a balanced model for everyday work; and Luna, a fast and affordable model. Terra has competitive performance to GPT-5.5 while being 2x cheaper and Luna brings strong capability at our lowest cost.

GPT-5.6 Sol launches with our most robust safety stack to date. We strengthened protections for higher-risk activity, sensitive cyber requests, and repeated misuse, and spent multiple weeks finding weaknesses, pressure-testing our system, and hardening it against real-world attacks.

We believe in broad access, and we plan to make GPT-5.6 Sol, Terra, and Luna generally available in the coming weeks. As part of our ongoing engagement with the U.S. government, we previewed our plans and the models' capabilities ahead of today's launch. At their request, we are starting with a limited preview for a small group of trusted partners whose participation has been shared with the government, before releasing more broadly. During this preview, we will continue testing and coordinating closely with partners as we work toward broader availability. We don't believe this kind of government access process should become the long-term default. It keeps the best tools from users, developers, enterprises, cyber defenders, and global partners who need them. We are taking this short-term step because we believe it is the strongest path to broader availability in the coming weeks, while we work with the Administration to develop the cyber Executive Order framework and a repeatable process for future model releases.

Capabilities

GPT-5.6 Sol is our strongest model yet. To give a preview of model performance, we share a set of evaluations highlighting improved agentic capabilities in coding, biology, and cybersecurity, with additional safety and preparedness evaluations available in our system card([opens in a new window](#)). We will share an expanded suite of evaluation results when we make the model broadly available.

With GPT-5.6, we're introducing a new `max` reasoning effort to give Sol the most time to reason deeply. Additionally, we're introducing a new `ultra` mode that goes beyond the capabilities of a single agent by leveraging subagents to accelerate complex work.

For coding workflows, GPT-5.6 Sol sets a new state of the art on **TerminalBench 2.1**, which tests command-line workflows requiring planning, iteration, and tool coordination.

TerminalBench 2.1

GPT-5.6 Sol also shows broad improvements in biology workflows. On **GeneBench v1**, which evaluates long-horizon genomics and quantitative-biology analyses, it achieves stronger results than GPT-5.5 while using fewer tokens.

GPT-5.6 Sol is our most capable model yet for cybersecurity. It shifts the performance-efficiency frontier for long-horizon security tasks including vulnerability research and exploitation. On **ExploitBench**², GPT-5.6 Sol is competitive with Mythos Preview using only ~1/3 of the output tokens. On **ExploitGym**³([opens in a new window](#)), a benchmark created by UC Berkeley researchers in collaboration with OpenAI and other frontier labs, GPT-5.6 Sol, Terra, and Luna models all demonstrate strong improvements in cyber capabilities as we increase reasoning.

Stronger cyber capabilities with stronger safeguards

We developed GPT-5.6 Sol, Terra and Luna with our most robust safeguards to date, with configurations matched to each model's capabilities. As the model becomes more capable, we design safeguards to increasingly hold up to real-world adversarial pressure while preserving access to legitimate work such as code review, vulnerability research, patch development, debugging, security education, and defensive testing. Our goal is to make prohibited offensive activity more difficult, uncertain, and detectable without unnecessarily limiting those beneficial uses. Based on our assessment of the model and safeguards, we expect substantial benefit for legitimate defensive work, while meaningfully constraining prohibited offensive use.

GPT-5.6 Sol is better at helping people find and fix vulnerabilities than reliably carrying out end-to-end attacks. As these capabilities continue to advance, our priority is to make sure they reach and benefit defenders, who can use these tools to find weaknesses, develop patches, and strengthen systems more broadly.

GPT-5.6 Sol does not cross the Cyber Critical threshold under our Preparedness Framework. In evaluations involving Chromium and Firefox, it identified bugs and exploitation primitives—the building blocks of an exploit—but did not autonomously produce a functional full-chain exploit under the conditions tested. Still, benchmark thresholds cannot capture every way a model may be used or combined with other tools. That uncertainty, along with the model's broader step change in capabilities, is why we are pairing the model's increased capabilities with stronger safeguards and a phased release. We share more details about our safeguards in the GPT-5.6 Preview system card([opens in a new window](#)).

A layered safeguard stack

No single safeguard is sufficient against determined or adaptive misuse. Across the GPT-5.6 preview, we use layered safeguards, with exact configurations varying across models, and pressure-test them for real-world attacks. These include protections trained into the model, real-time checks during generation, account-level signals, differentiated access, monitoring, enforcement, and continued testing.

GPT-5.6 is trained to refuse prohibited cyber assistance, including when users attempt to disguise their intent or jailbreak the model. These model-level safeguards establish the first boundary around what the model should and should not help with.

Real-time cyber and biology misuse classifiers provide another layer by evaluating output as it is generated. For higher risk cases, if they detect a potential violation, the generation may be paused while a larger reasoning model reviews the conversation and its context. If the output is assessed as disallowed, it is withheld before it reaches the user.

Flagged activity can also trigger account-level review across relevant conversations and risk signals, consistent with our terms and policies around content retention and review. Looking beyond a single conversation helps our systems distinguish persistent malicious behavior from legitimate dual-use security work, where similar technical concepts may appear in very different contexts.

Together, these layers make the overall approach more robust than any one safeguard on its own. Model behavior reduces the likelihood of harmful responses, real-time systems can intervene during generation, account-level review can identify broader patterns, and differentiated access preserves important defensive work without making the most sensitive capabilities broadly available by default.

Especially during the preview, users may encounter safeguards that block or refuse some requests. Other requests may take longer because generation is paused for additional review. Safeguards may occasionally intervene on legitimate work, particularly in dual-use areas where defensive and offensive activity can initially look similar.

That is part of what the preview is designed to test. We want to understand not only whether the safeguards constrain misuse, but whether legitimate users can still complete normal work reliably and efficiently. Feedback during the preview will help us reduce unnecessary blocks and delays, improve how the safeguards interpret context, and create a smoother experience before wider release.

We are also working with enterprise customers on longer-term approaches—including privacy-preserving detection, customer-operated safety controls, and access calibrated to the risk of a customer, user, or workload—to advance safety while supporting enterprise privacy requirements.

Improving robustness with automated red-teaming

Safeguards also need to remain effective when attackers adapt their tactics. A protection that works only on a fixed set of known attacks is not robust enough for a frontier model.

That's why we are applying more intelligence and compute than ever before to safety, using our own models to find weaknesses and improve safeguards faster. We dedicated over 700,000 A100-equivalent GPU hours to automated red teaming aimed at finding universal jailbreaks: attacks that can work across many prompts or contexts, not just one narrow setting. Focusing on these harder, more general attacks let us test the safeguards beyond a fixed set of known failures. It also lets us explore far more attack patterns than human testing alone could cover, identify failure patterns earlier, and shorten the path from finding a weakness to addressing it.

In addition to automated red-teaming, we worked with third-party testers to conduct extensive human expert red teaming, which will continue in the preview period. Human red-teaming complements the automated work by testing safeguards against creative experts trying to misuse the model in ways our systems might not anticipate.

No evaluation can represent every product configuration, multi-step attack, or real-world workflow. We therefore maintain a rapid-response process to reproduce, assess, prioritize, and remediate newly discovered jailbreaks, then add them to our ongoing evaluations so we can test against similar failures in the future.

Availability and pricing

During the preview, GPT-5.6 models will initially be available through the API and Codex to a select group of trusted partners and organizations. We plan to make them more broadly available to people using ChatGPT, Codex, and the API soon.

In this new naming system introduced with GPT-5.6, the number identifies a model's generation, while Sol, Terra, and Luna identify durable capability tiers that can advance on their own cadence. Together, the family gives people and developers clearer choices across intelligence, speed, and cost.

GPT-5.6 is priced per 1M tokens across three model sizes: Sol is \$5 input / \$30 output; Terra is \$2.50 input / \$15 output; and Luna is \$1 input / \$6 output. GPT-5.6 also introduces more predictable prompt caching, including support for explicit cache breakpoints and a 30-minute minimum cache life. For GPT-5.6 and later models, cache writes are billed at 1.25x the model's uncached input rate, while cache reads continue to receive the 90% cached-input discount.

We're also launching GPT-5.6 Sol on Cerebras at up to 750 tokens per second in July, bringing frontier intelligence to customers at unprecedented speed. Access will initially be limited to select customers as we expand capacity.

We're excited to continue learning from this preview period, and to bring GPT-5.6 Sol, Terra and Luna to more people soon.

* * *

1. We estimate latency and API cost by looking at the production behavior of our models, and simulating offline. These estimates account for tool call details, sampled tokens, and input tokens. Real-world results may vary substantially, and depend on many factors not captured in our simulation. We simulate latency at fast API speeds, and cost at regular API pricing.
2. All models are evaluated using the ExploitBench API harness with 5 seeds and reasoning continuity.
3. We ran ExploitGym on our alpha API, which outputs responses faster than our public API, and then rescaled to match our public API. When rescaling latencies to the speeds expected for our public API, this causes some estimated latencies to exceed the 2h and 6h hour time limits, despite being correctly obeyed in the evaluation run. To get faster speeds for time-sensitive work, we offer priority processing in the API and fast mode in Codex.
4. Models without reported output tokens, latency or cost are plotted as horizontal dotted lines.

* 2026

Author

OpenAI

Keep reading

[View all](#)

[Introducing the Rosalind Academic Plan Research](#) Oct 15, 2026

[Helping teens learn, plan, and shape the future of AI Product](#) Oct 7, 2026

[Video 2](#)

[GPT-6 and Intelligent UI for everyone Product](#) Oct 7, 2026

[Research](#) * [Research Index](#) * [Research Overview](#) * [Economic Research](#)

[Latest Advancements](#) * [GPT-6.1 Sol](#) * [GPT-6 Astra](#) * [GPT-5.6](#) * [GPT-5.5](#)

[Safety](#) * [Safety Approach](#) * [Deployment Safety](#)(opens in a new window) * [Security & Privacy](#) * [Trust & Transparency](#)

[Products](#) * [ChatGPT](#)(opens in a new window) * [ChatGPT Business](#)(opens in a new window) * [ChatGPT Enterprise](#)(opens in a new window) * [ChatGPT for Education](#)(opens in a new window) * [Codex](#) * [Dots](#)(opens in a new window) * [Release Notes](#)

[API Platform](#) * [Overview](#) * [API Log In](#)(opens in a new window) * [Docs](#)(opens in a new window)

[Business](#) * [Overview](#) * [Solutions](#) * [Resources](#) * [Plugins](#) * [Customer Stories](#) * [Partner Network](#) * [Contact Sales](#)

[Developers](#) * [Apps SDK](#)(opens in a new window) * [Open Models](#) * [Docs](#)(opens in a new window) * [Resources](#)(opens in a new window) * [Developer Forum](#)(opens in a new window)

[Company](#) * [About Us](#) * [Our Charter](#) * [Careers](#) * [News](#)

[Support](#) * [Help Center](#)(opens in a new window) * [OpenAI Status](#)(opens in a new window)

[More](#) * [Stories](#) * [Academy](#) * [Supply Co.](#) * [Livestreams](#) * [Podcast](#) * [RSS](#)

[Terms & Policies](#) * [Terms of Use](#) * [Privacy Policy](#) * [Other Policies](#)

(opens in a new window)(opens in a new window)(opens in a new window)(opens in a new window)(opens in a new window)(opens in a new window)(opens in a new window)

OpenAI © 2015–2026 [Your privacy choices](#)

English United States